GASCN: Graph Attention Shape Completion Network

Haojie Huang   Ziyi Yang   Robert Platt
Northeastern University
360 Huntington Avenue, Boston, MA 02115, United States
{huang.haoj; yang.ziyi2} @northeastern.edu; rplatt@ccs.neu.edu

Abstract

Shape completion, the problem of inferring the complete geometry of an object given a partial point cloud, is an important problem in robotics and computer vision. This paper proposes the Graph Attention Shape Completion Network (GASCN), a novel neural network model that solves this problem. This model combines a graph-based model for encoding local point cloud information with an MLP-based architecture for encoding global information. For each completed point, our model infers the normal and extent of the local surface patch which is used to produce dense yet precise shape completions. We report experiments that demonstrate that GASCN outperforms standard shape completion methods on a standard benchmark drawn from the Shapenet dataset.

1. Introduction

It is often desirable to be able to infer the geometry of objects in a scene based on a small number of sensor measurements. This is known as the *shape completion* problem where the system infers the complete geometry of an object based on a partially observed point cloud, such as that produced by a depth sensor. In particular, shape completion could be useful in the context of robotic manipulation where a robotic system could use the inferred geometry of the objects in a scene to create geometric plans that solve desired tasks. Compared to other approaches to scene reconstruction that estimate the pose of specific objects, a critical advantage of shape completion methods is that they can be used to infer the geometry of novel objects, not just a closed set of previously modeled objects.

There are two main types of approaches to shape completion: one based on the point cloud representation and the second based on voxel grids or signed distance functions. Point clouds are a convenient choice because they can be created easily from the output of depth sensors or LIDAR and several methods use point cloud representations [36, 35, 8, 17]. However, these methods are all limited by the fact that they cannot use 3D convolutions because the point cloud representation does not explicitly encode information about the local connectivity of each point. An alternative is to create shape completion algorithms based on a voxel grid or signed distance function (SDF) representation [22, 12, 24]. Here, we can use 3D convolutions to obtain higher shape completion accuracy, as in GRNet [32]. However, the memory requirements for voxel grids or SDFs are cubic in the resolution of the grid and therefore it severely limits the resolution at which the method can be applied.

In contrast to the methods above, this paper proposes GASCN, a shape completion approach that operates on a graph rather than a point cloud or voxel grid. Since converting a point cloud into a mesh can itself be challenging, we use a graph attention approach [26] where each point in the cloud is associated with a unique graph that connects it with its neighbors. Another key element of our approach is a densification step that converts a coarse reconstruction into a fine reconstruction by deforming a fine 2D grid to each coarse point via a learned surface normal and variance. The novel features of our approach relative to prior shape completion work are the following:

1. Our encoder utilizes the graph attention layer to encode local context information for further combination with the global structure information whereas prior work on shape completion typically uses a PointNet [20] encoder or 3D convolutions on voxels [31]. This enables us to take point cloud input while still encoding using graph convolutions.

2. Our decoder contains a novel densification procedure that maps a dense planar grid to a local region around each point in a coarse shape completion. This enables us to produce a dense completion while keeping the network complexity manageable.

We compare our model with several other recent approaches from the literature [35, 36, 25] and find that it outperforms on a standard shape completion baseline task.

2. Related Work

The problem of inferring 3D structures from RGB and depth camera images has been studied for decades. A classic but powerful non-learning method uses the Iterative Closest Point (ICP) [1] algorithm to align a template model from a large dataset with the observed point cloud. BEOs [3] constructs a new representation space with VBPCA [2] basis, where the best corresponding point to the partial input could be found. Some other works [21, 10, 16] deform the template model to synthesize shapes that are more consistent with the observations. These methods are often sensitive to noise and require large datasets as well as expensive computation time.

Recent learning-based approaches complete the 3D structure with a parameterized model, often a deep neural network, directly from the observation. It allows a fast inference and great generalization. Some datasets like ShapeNet [5] and YCB [4] enable learning on sufficient examples to generate visually compelling shapes. The most common structure for a shape completion network consists of an encoder that maps the partial observation to a latent vector and a decoder that produces the complete 3D shape given the latent vector. The learning method differs based on the different representation methods of a 3D object.

Voxel-based methods [24, 31, 12] voxelize point clouds into binary voxels, where 3D convolutional neural networks can be applied. Recently, GRnet [32] proposed a novel gridding mechanism to encode the neighbors' information into the vertices of each voxel. These methods are often computation expensive and limited to quantization effects due to the voxel representation.

PointNet [20] is a pioneer in using the MLP-based approach to learn point clouds. It combined pointwise multilayer perceptrons with a max-pooling aggregation function to achieve invariance to permutation and robustness to perturbation. After PointNet, FoldingNet [35] proposed a folding-based decoder that deforms a 2D grid onto the 3D surface of a point cloud. Following previous works, PCN [36] proposed the coarse-to-fine procedure to densify the point cloud. AtlasNet [8] and MSN [17] generated a collection of parametric surface elements based on FoldingNet's decoder. Besides, some work [15, 30] combined MLPs with GANs [7], and some [22] integrated it with reinforcement learning. However, these approaches just aggregated the global information from the whole point cloud without focusing on the local information of each point's neighbors, and neither did they take the various local properties of different points into account during the generation process.

Compared with MLP-based approaches, graph-based methods could extract local information of each point. Most recent graph-based convolution methods [14, 11, 26] operated on groups of spatially close neighbors. ECC [23] is the first one to apply graph convolutions to point cloud classification. After that, DGCNN [29] proposed a more generalized edge convolution method with dynamic graph updates to encode the local neighborhood information for point clouds. Inspired by DGCNN, DCG [27] introduced this method to point cloud completion following a coarse-to-fine fashion. However, the edge convolution is computationally expensive in large graphs. After Graph Attention Network [26], GAC [28] inherited its ideas and applied attention mechanism to the point cloud segmentation area. To the best of our knowledge, we are the first to directly combine the graph-level local information with the global structure information in the encoder as well as utilize surface normals to incorporate pointwise local properties for shape completion.

3. Method

Figure 1 below is an overview of our network, which follows an encoder-decoder architecture. It takes a partial point cloud with a variable number of points as input, p_{input} and generates a complete dense point cloud as output, p_{final}.

The encoder (Figure 2) uses a graph attention network (GAT) [26] to reason about the local geometry of the input cloud and MLPs to reason about global geometry.

GAT [26] is used to learn the graph-level local information and MLPs are used to encode the global information. Following a coarse-to-dense procedure, our decoder first decodes a sparse point cloud P_c and then samples each point's neighbors to densify P_c. To make the densifying process flexible and adaptive to different local geometries, we decode surface normals and the scale factor σ by our normal decoder and sigma decoder to adjust the orientation and the size of each point's neighborhood.

3.1. Encoder

Since each region of a point cloud could have different geometries, all the global information without the local information would result in averaging over all point features and thus undermining the completions. In order to solve it, our encoder is built on the graph representation of partial point clouds. We encode each node feature with its neighbors' information, its mapped Cartesian feature, and two different types of global vectors to maximize the utilization of local information and global information.

3.1.1 Build Graph for Point Cloud

Given a point cloud consisting of m points, denoted as $P = \{p_1, p_2, ..., p_m\}$ with $p_i \in \mathbb{R}^3$, we form a directed graph $G = (V, E)$ representing the point cloud structure, where $V = \{1, ..., m\}$ and $E \subseteq V \times V$. We define $H = \{h_1, h_2, ..., h_m\}$ with $h_i \in \mathbb{R}^D$ as a set of node
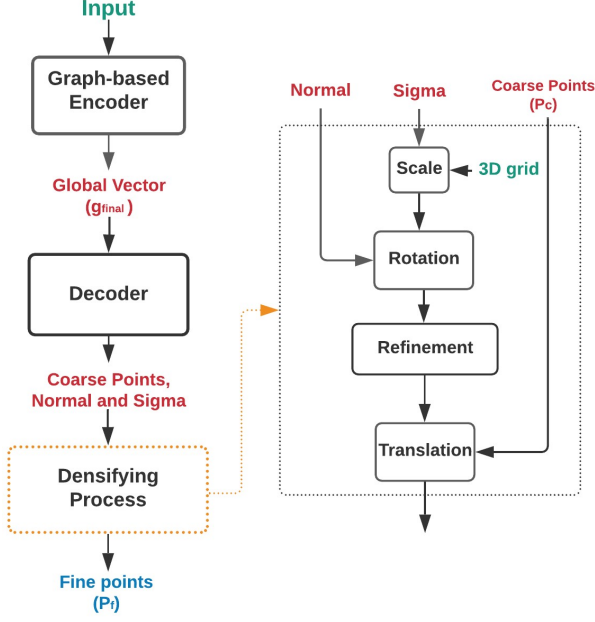
Figure 1: The architecture of GASCN. It has a graph-based encoder that integrates each point's local information with the point cloud's global structure information. The decoding process utilizes surface normals and coarse points to rotate and translate adaptive 3D grids to densify coarse point clouds.

features corresponding to each point, where D is the dimension of the node features. Each node feature is initialized with $h_i = (x_i, y_i, z_i)$, where (x_i, y_i, z_i) is the Cartesian coordinates of point p_i. Moreover, we build subgraphs $\{G_i\}_{i=1,2,...,m} \subseteq G$ for all nodes as k-nearest neighbor [19] graphs from node features H by connecting the corresponding point p_i with its k-nearest spatial neighbors. Besides, the subgraphs include self-loops to preserve each node's own information.

As a result, the whole graph for the point cloud comprises m directed subgraphs and each point is assigned as the destination node of a subgraph. Each of them seems like a patch on the surface of a 3D object, similar to the kernel window in 2D convolution.

3.1.2 Local Convolution: GAT

Although points in the point cloud are independent, there must exist certain hidden relationships to maintain the local surface geometry. GAT [26] provides an attentive mechanism to learn it. Given the graph G and point features H of a partial point cloud, we apply the GAT layer to aggregate each node's neighborhood information.

First, we conduct a linear transformation for the embedding feature h_i of the node i with a shared learnable weight
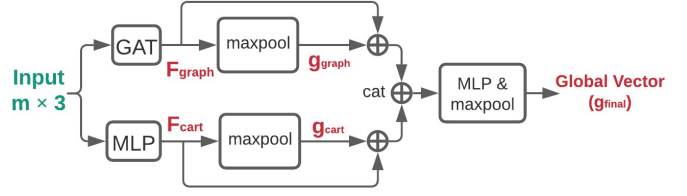

Figure 2: Graph-based Encoder of GASCN model. It encodes each node feature with its neighbors' information F_{graph} by graph attention layer, its mapped Cartesian feature F_{cart} by point-wise MLP, and two types of global vectors, g_{graph} and g_{cart}, extracted from the max-pooling operation.

matrix W and then compute a pair-wise unnormalized attention score e_{ij} with a learnable vector a between the two neighbors i and j ($||$ is the concatenation operation).

$$z_i = W h_i \tag{1}$$

$$e_{ij} = LeakyReLU(a^T(z_i||z_j)) \tag{2}$$

After that, our network normalizes the attention score on each node's neighbors and aggregates the embeddings from neighbors with the normalized attention score, and applies a nonlinear activation function.

$$\alpha_{ij} = \frac{exp(e_{ij})}{\sum_{k \in N(i)} exp(e_{ik})} \tag{3}$$

$$h_i^{new} = LeakyReLU(\sum_{j \in N(i)} \alpha_{ij} z_j) \tag{4}$$

The final result is denoted as F_{graph} which is an $m \times d$ matrix, where m is the number of points and d is the feature dimension. The process above maps each node feature from a 3-dimension xyz coordinate to a d-dimension vector that contains the information of its neighbors through the graph attention layer.

3.1.3 Global Information

The global structure information is critical in point cloud generation, which could be used to distinguish different objects. We extract two different types of global information, the global graph vector g_{graph} and the global Cartesian vector g_{cart}. The global graph vector g_{graph} is given by

$$g_{graph} = maxpool(F_{graph}) \tag{5}$$

where max-pooling operation concatenates the maximum value among each feature dimension of F_{graph} to construct the permutation-invariant global vector g_{graph}.

We define the pointwise MLPs in PointNet [20] as $M_\theta : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$. It maps each point feature from d_1-dimension

into a new d_2-dimension space using a nonlinear function M with the shared learnable weights θ. The second global vector is aggregated by:

$$F_{cart} = M_{\theta_1}(H) \tag{6}$$

$$g_{cart} = maxpool(F_{cart}) \tag{7}$$

After that, two different feature matrices, F_{graph} and F_{cart}, are concatenated into an augmented feature matrix $\bar{F}$, while two global vectors g_{graph} and g_{cart} are concatenated together to produce the augmented global vector $\bar{g}$. We concatenate $\bar{g}$ to each row of $\bar{F}$ to get our final feature matrix F_{final}. So far, each node feature is encoded with information of its neighbors, its mapped Cartesian feature, and two types of global vectors. Finally, we extract the final latent vector g_{final} by:

$$g_{final} = maxpool(M_{\theta_2}(F_{final})) \tag{8}$$

3.2. Decoder

Regarding the various local properties of different points during the generation process, our decoder first generates coarse points, each coarse point's surface normal, and its corresponding scale factor σ, as shown in Figure 3. It then samples coarse points' neighbors with adaptive 3D grid points and utilizes both surface normals and coarse points to determine the rotation and the translation, respectively.

3.2.1 Coarse Point Cloud and Surface Normal

The coarse output P_c is a preliminary generated complete shape with sparse points. Following previous work [36], we use three fully connected layers with nonlinear activation functions to decode the g_{final} to P_c, as shown at the top of Figure 3.

In order to densify completions, some pointwise attributes are needed to generate neighbors. Surface normal is a good option which contains the local orientation of a surface. To decode surface normals, we first build a graph by connecting each coarse point with its nearest neighbors and apply one graph convolution layer [14] to aggregate the neighbor information into C_{local}. In the meantime, the global vector g_{final} and coarse points P_c are mapped respectively by two different MLPs into G_{map} and C_{map}. The surface normal is generated by

$$N_{imd} = M_{\phi_1}(C_{local} \,||\, C_{map} \,||\, G_{map}) \tag{9}$$

$$Normal = \frac{M_{\phi_2}(N_{imd})}{||M_{\phi_2}(N_{imd})||} \tag{10}$$

where the N_{imd} is the normal internal embedding. The whole process above is shown in Figure 3.
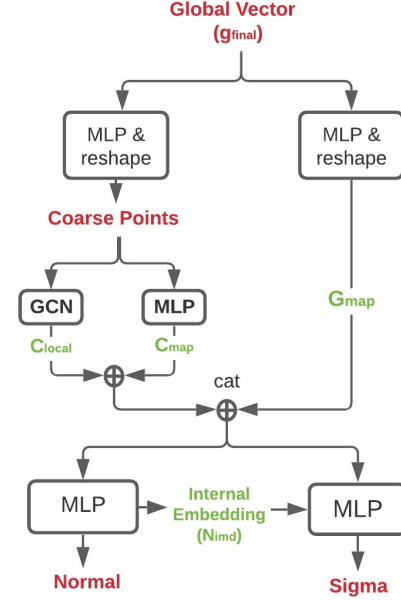

Figure 3: Decoder Architecture. The coarse points decoder consists of three fully connected layers. The normal decoder takes as inputs the mapped coarse point coordinate C_{map}, its neighbors' information C_{local}, and the mapped global vector G_{map}. The sigma decoder takes as inputs the normal internal embedding N_{imd}, C_{local}, C_{map}, and G_{map}, and then applies three pointwise MLP layers to output pointwise-attentive sigmas.

3.2.2 Calculate Neighbors by Surface Normals in Coarse Point Cloud

Given a point $\mathbf{p} = (p_x, p_y, p_z)$ and its normal $\mathbf{n} = (n_x, n_y, n_z)$, how could we sample its neighbor points? Neither sampling X, Y, Z randomly [17] from a distribution that would make the output look noisy and the model hard to train, nor utilize a fixed grid which is inflexible; we instead mesh an adaptive grid to obtain a set of 3D Cartesian coordinates for each coarse point. Our procedure has four steps: sampling, scaling, rotation, and translation.

1. Obtain neighbors by meshing grid:

$$(X, Y) = Meshgrid(l, l); \; Z = \mathbf{0} \tag{11}$$

(X, Y, Z) is a set of vertices of a square grid of size l, centered at the origin with surface normal oriented along Z-axis, forming a plane in 3D space.

2. Scale X, Y with σ:

$$x = \sigma \cdot X; y = \sigma \cdot Y \tag{12}$$

Each point has a specific scale factor σ that is learned from the neural network, which we will discuss in Sec-

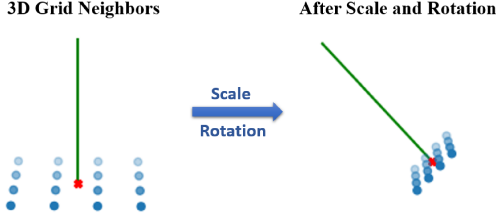tion 3.2.3. After the scaling, different points are associated with planes of different sizes accordingly.

Figure 4: After the scale and rotation, we align the initial 3D-grid neighbors' surface normal along the decoded $Normal$. This process enables our generated grid points to match local surfaces with different orientations and geometries. (Blue dots are neighbors, the green line indicates the normal direction, and the red dot is the center.)

3. Rotate the samples to make their surface normals align along the surface normals we generated, as shown in Figure 4:

$$k = z_{axis} \times n \qquad (13)$$

$$\theta = \arccos(n_z) \qquad (14)$$

where the k is the rotation axis, z_{axis} is the unit vector along Z-axis, $\times$ is the cross product, θ is the rotation angle. If k is a zero vector, no rotation is needed. We use Rodrigues' Rotation Formula to transform the axis-angle representation to a rotation matrix R.

4. Translate the center of the samples to point $\mathbf{p}$. The samples are centered at the origin $(0,0,0)$ previously, and we add the generated point's coordinate (p_x, p_y, p_z) to each sample's coordinate to complete the translation.

After applying the operations to each point of P_c, we fuse each point's local surface information with its coordinate to densify the point cloud.

3.2.3 Scale Factor σ

σ is the scale factor learned from our sigma network shown in Figure 3 to scale the samples. For each point in P_c, there is a σ corresponding to it. It makes the neighbor generation be flexible and attentive. Using a table as an instance, the corner points or edge points may have smaller σ, while points on a flat plane should have a larger σ since they have more confidence in sampling neighbors with a wider grid. In our work, σ is pointwise attentive and learned by

$$\log \sigma = M_{\phi_\sigma}(N_{imd} \,||\, C_{local} \,||\, C_{map} \,||\, G_{map}) \qquad (15)$$

3.2.4 Deformation and Refinement

So far, our network generates a dense set of points, P_d. In order to deform the 3D-grid points and ensure robust performance, each point is concatenated with the corresponding coarse point, which serves as its positional embedding. We then map the feature to a high-dimension space and project them back to 3D space by three points-wise MLPs and Leaky-Relu activation function [33]. We intuitively denote the improved dense point clouds as fine point clouds, P_f, which marks the end of our forward propagation.

3.3. Loss

Since the point cloud is an unordered set, we use the permutation-invariance loss, Chamfer Distance (CD) [6], to measure the distance between our output and the ground truth. CD (16) calculates the average closest point distance between two sets of point clouds. It has two terms to force the output to be closed to as well as cover the ground truth.

$$
\begin{aligned}
CD_{loss}(S_1, S_2) = &\frac{1}{|S_1|} \sum_{x \in S_1} \min_{y \in S_2} ||x - y||_2 \\
&+ \frac{1}{|S_2|} \sum_{y \in S_2} \min_{x \in S_1} ||y - x||_2
\end{aligned}
\qquad (16)
$$

Our network generates two outputs: Coarse points P_c, fine points P_f, and each output corresponds to a CD loss with respect to the ground truth.

$$Loss = CD_{loss}(P_c, P_{gt}) + \alpha \, CD_{loss}(P_f, P_{gt}) \qquad (17)$$

where α is a hyperparameter and P_{gt} are ground truth points.

4. Experiments

This section describes our dataset and the training process first, and then compares our methods with several strong baselines. Afterward, we present the ablation study results, followed by experiment on real-sensor data. Following previous works [36, 25, 32], Chamfer Distance serves as our quantitative evaluation metric for baseline comparison and ablation study.

4.1. Datasets and Implementation Details

ShapeNet PCN. The ShapeNet dataset [5] for point cloud completion is derived from PCN [36], consisting of 30,974 3D models from 8 categories. The ground truth point clouds containing 16,384 points are uniformly sampled on mesh surfaces. The partial point clouds are generated by randomly sampling 8 camera poses to capture 8 depth images for each CAD model and backing-projecting the 2.5D depth images to 3D, and thus own a variant number of points.

Model	Airplane	Cabinet	Car	Chair	Lamp	Sofa	Table	Watercraft	Overall
FoldingNet	10.74	19.31	21.03	21.89	19.93	19.82	22.55	18.31	19.20
GRnet	9.16	17.98	14.96	12.34	10.66	14.56	13.84	11.05	13.06
TopNet	7.12	14.54	13.19	13.65	11.51	12.80	12.88	10.13	12.05
PCN	6.51	14.43	11.81	10.90	10.32	10.92	14.58	8.96	11.05
Ours	**4.71**	**11.69**	**9.67**	**8.44**	**6.91**	**8.31**	**10.02**	**6.69**	**8.31**

Table 1: Point completion results comparison on ShapeNet using Chamfer Distance (CD) with L2 norm computed on 16,384 points and multiplied by 10^3. The best results are highlighted in bold.
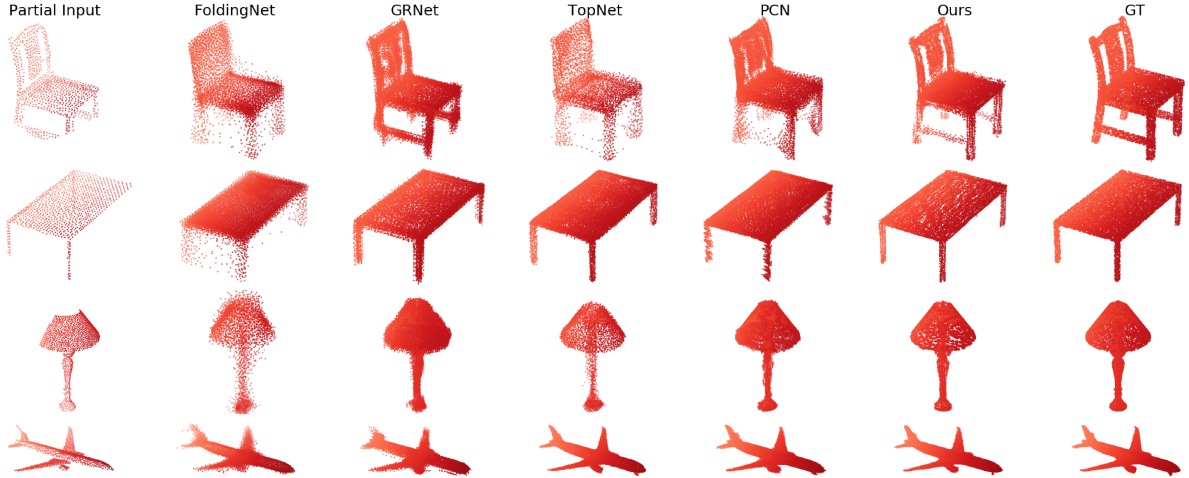

Figure 5: Visual Results Comparison. A partial point cloud is given and our method generates better complete point clouds.

ShapeNet Bottle. We built a small dataset for ablation study and our future work in manipulator grasping following ShapeNet PCN. It contains 498 bottle-shaped objects such as water bottles, jars, spray bottles, etc. Each object is also used to generate eight partial observations with a flexible number of points and one ground truth point cloud with 9,216 points. 80% of the models are used for training, 10% for evaluation and 10% for testing.

YCB Dataset. It consists of objects of daily life with different shapes, sizes, and textures. The dataset provides mesh models and RGB-D scans of the objects.

We select the following setting based on our implementations and dataset. During constructing the graph representation of the point cloud, each point is connected with its 20 neighbors. KDTree [18] was used for fast retrieval of nearest neighbors. We first generate a coarse point cloud with 1,024 points and link each point with its 5 nearest neighbors. A 4×4 grid initialized with l equals to 0.1 is used to sample 16 neighbors for each coarse point. As a result, our model consumes an input with a variable number of points and generates the coarse output with 1,024 points and fine output with 16,384 points. We trained the networks for 200 epochs with a decaying learning rate initialized at 10^{-4} and Adam optimizer [13], with a batch size of 32 on Four

Method	FN	GRnet	TopNet	PCN	Ours
# Params	2.40M	76.7M	3.27M	6.85M	9.60M

Table 2: Number of trainable model parameters

NVIDIA GeForce RTX 2080Ti GPUs. It takes about 78 hours to complete 100 epochs. The average inference time for one instance is 0.075 seconds on a single GPU.

4.2. Baselines and Performance Comparison

Here, we compare our model against four baselines that work on point cloud completion.

1. **FoldingNet:** It encodes each point's neighbor information by flattening the covariance matrix and proposed a folding-based decoder that deforms a 2D fixed grid onto the 3D surface of a point cloud [35].

2. **TopNet:** It proposes a decoder following a hierarchical rooted tree structure to generate points [25].

3. **GRnet:** A recent method that uses 3D voxels as intermediate representations for point clouds [32].

4. **PCN:** The representative of point completion methods following the coarse-to-fine fashion [36].

Quantitative results in Table 1 and Table 2 indicate our method could achieve better performance over all eight categories while still maintain a reasonable number of parameters. As shown in Figure 5, our completion results could recover a more detailed geometry than others. It proves the advantage of our proposed method that extracts both the local information and the global structure information from the observations and the strength of generating various local geometries.

4.3. Ablation Study

In order to investigate our encoder and decoder's performance separately, we train four different neural networks for the ablation study on the ShapeNet Bottle. A 3×3 grid is used to generate 9,216 points. We evaluate the performance on the validation dataset every 20 epochs during training.

1. **Model A**: Removing surface normal and adaptive 3D grid from our decoder (same as PCN's decoder).
2. **Model B**: Removing GAT from our encoder (same as PCN's encoder).
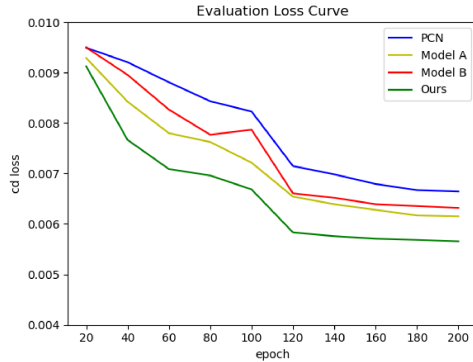

Figure 6: CD loss during evaluation. The vertical axis is the CD loss, and the horizontal axis is the epoch number. We evaluate the performance every 20 epochs on the validation dataset.

As shown in Figure 6, both our encoder and decoder converge fast and outperform PCN's encoder and decoder. Without GAT in the encoder, the performance is decreased by 8%; without surface normals and scale factor σ, the performance is decreased by 10%. It further indicates our encoder's benefit, which combines local information with global information, and our decoder's power, which utilizes surface normals and adaptive grid sizes.

We also investigated using multiple GAT layers to encode local information, while the test results show no improvement but the deterioration of the performance. We speculate the main reason is that two graph attention layers allow each point to communicate with more neighbors,
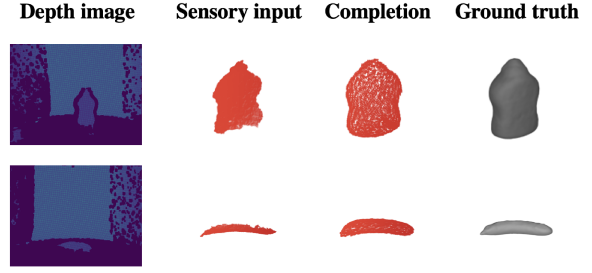

Figure 7: Experimental results on YCB objects, showing completions for a mustard bottle and a banana, respectively. We extract the point clouds (second column) from the raw depth images (first column), and feed them into the model to generate the completions (third column). The ground truth meshes are shown in the last column.

which will lead it to lose the property of local convolution, and hence be difficult to learn efficiently with the attention mechanism.

4.4. Experiment on Sensory Data

We test our model on YCB objects with a reasonable reconstructed mesh (70 in total), as shown in Figure 7. We first build a small simulation dataset using the same strategy mentioned in ShapeNet Bottle. After training our model for 150 epochs, we test it on sensory point clouds extracted from raw registered depth images. Our model achieves an average chamfer distance of 9.9834×10^{-3} on depth images of the first camera view of each object from the YCB dataset. It is worth noting that the real sensor data is nosier and occasionally suffers from missing geometry features compared to the simulated data.

5. Discussion

5.1. Nearest Neighbor Distance for Each Point

Since Chamfer Distance only provides the average performance for the whole point cloud [34], to analyze our results in detail, we illustrate each point's nearest neighbor distance to the ground truth in Figure 8. The generated points are rendered by their closest distances to the ground truth. As it suggests, each point maintains a relatively small distance to the closest point in the ground truth.

5.2. Generalization Ability

When testing on the ShapeNet Bottle dataset, we found some interesting results. Figure 9 shows our generation process for a shampoo bottle and a coke can with a straw. There is no straw-like object but some shampoo bottles in the training dataset, and our network could still reason out the geometry of the straw, though it resembles the shampoo
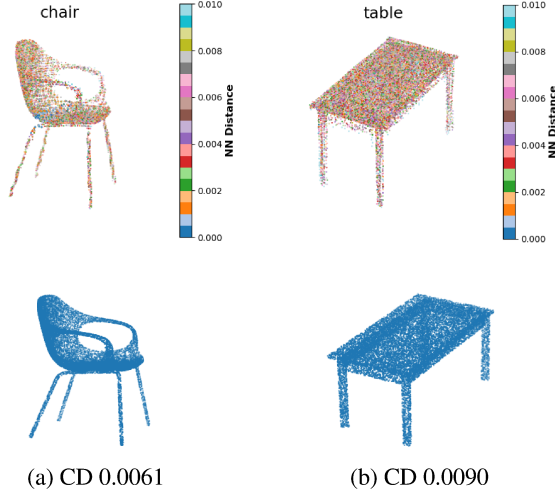
(a) CD 0.0061   (b) CD 0.0090

Figure 8: Nearest Neighbor Distance for Each Point. The top row shows our generated outputs, and the bottom row shows the ground truth. The points are rendered by their closest distances to the ground truth.

header to some degree. This phenomenon shows that our model could predict a reasonable shape for the unseen part of the input point cloud, which indicates the generalization ability of our proposed method.
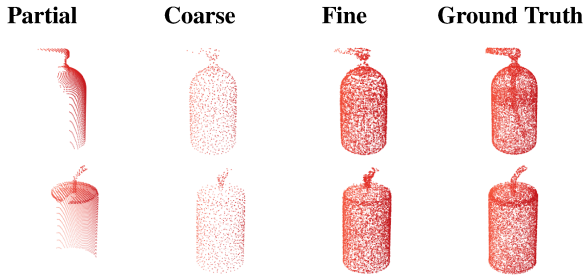

Figure 9: Illustration of Generalization Ability of GASCN: first, we take a partial point cloud as input and output coarse point clouds, and then do the coarse-to-fine refinement. The top row and the bottom row show the generation process of a shampoo bottle and a coke can with a straw, respectively. Though there is no straw in our training data, our model generates a similar geometry for the unseen part of the partial observation during testing.

5.3. Point Cloud Registration

Many tasks in robotics could benefit from dense complete point clouds. We illustrate its advantage on the point cloud registration. Specifically, we ran ICP [1] for two different partial observations from one object and compared its results with their corresponding complete fine points. Our experiment indicated that our generated complete points

(a) MSE: 0.2915   (b) MSE: 0.0657

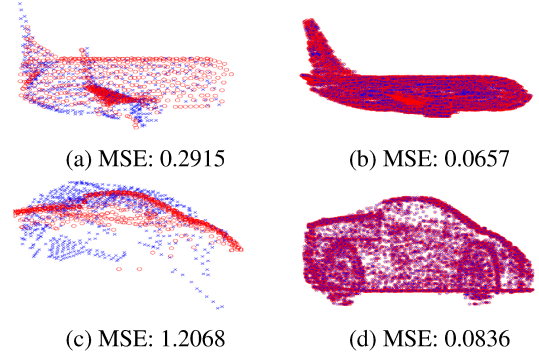(c) MSE: 1.2068   (d) MSE: 0.0836

Figure 10: Improvement on Point Cloud Registration. The left column shows registrations for partial observations, and the right column shows registration for their corresponding complete fine points.

have fewer errors since they have significant overlaps compared to the partial observations. It is worth noting that our model could enormously increase the utilization of observations in many downstream tasks, especially where both the completions and surface normals are needed.

6. Conclusion

In this work, we studied how to infer the complete 3D geometry from an incomplete observation. Our motivation is to combine local information with global information in our encoder and take the pointwise local property into account during our generation. To achieve the target, we utilized the graph attention layer for encoding local context information and utilized MLPs and max-pooling to aggregate the global structure information. To decode the latent vector properly, we inversed the process of calculating surface normal from points and utilized adaptive 3D grids to obtain neighbors. In our future work, we will investigate combining shape completion with grasp pose detection [9] in robotics.

Acknowledgement

References

[1] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. International Society for Optics and Photonics, 1992.

[2] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.

[3] Benjamin Burchfiel and George Dimitri Konidaris. Bayesian eigenobjects: A unified framework for 3d robot perception.

In *Robotics: Science and Systems*, volume 7, page 2017, 2017.

[4] Berk Calli, Arjun Singh, James Bruce, Aaron Walsman, Kurt Konolige, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research*, 36(3):261–268, 2017.

[5] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.

[6] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017.

[7] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.

[8] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 216–224, 2018.

[9] Marcus Gualtieri, Andreas Ten Pas, Kate Saenko, and Robert Platt. High precision grasp pose detection in dense clutter. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 598–605. IEEE, 2016.

[10] Saurabh Gupta, Pablo Arbelaez, Ross Girshick, and Jitendra Malik. Aligning 3d models to rgb-d images of cluttered scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[11] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. *arXiv preprint arXiv:1706.02216*, 2017.

[12] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *Proceedings of the IEEE international conference on computer vision*, pages 85–93, 2017.

[13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[15] Chun-Liang Li, Manzil Zaheer, Yang Zhang, Barnabas Poczos, and Ruslan Salakhutdinov. Point cloud gan. *arXiv preprint arXiv:1810.05795*, 2018.

[16] Dongping Li, Tianjia Shao, Hongzhi Wu, and Kun Zhou. Shape completion from a single rgbd image. *IEEE transactions on visualization and computer graphics*, 23(7):1809–1822, 2016.

[17] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11596–11603, 2020.

[18] Marius Muja and David G Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. *VISAPP (1)*, 2(331-340):2, 2009.

[19] Leif E Peterson. K-nearest neighbor. *Scholarpedia*, 4(2):1883, 2009.

[20] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

[21] Jason Rock, Tanmay Gupta, Justin Thorsen, JunYoung Gwak, Daeyun Shin, and Derek Hoiem. Completing 3d object shape from one depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2484–2493, 2015.

[22] Muhammad Sarmad, Hyunjoo Jenny Lee, and Young Min Kim. Rl-gan-net: A reinforcement learning agent controlled gan network for real-time point cloud shape completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5898–5907, 2019.

[23] Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3693–3702, 2017.

[24] David Stutz and Andreas Geiger. Learning 3d shape completion from laser scan data with weak supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1955–1964, 2018.

[25] Lyne P. Tchapmi, Vineet Kosaraju, Hamid Rezatofighi, Ian Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[26] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

[27] Kaiqi Wang, Ke Chen, and Kui Jia. Deep cascade generation on point sets. In *IJCAI*, volume 2, page 4, 2019.

[28] Lei Wang, Yuchun Huang, Yaolin Hou, Shenman Zhang, and Jie Shan. Graph attention convolution for point cloud semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10296–10305, 2019.

[29] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019.

[30] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *arXiv preprint arXiv:1610.07584*, 2016.

[31] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.

[32] Haozhe Xie, Hongxun Yao, Shangchen Zhou, Jiageng Mao, Shengping Zhang, and Wenxiu Sun. Grnet: gridding residual network for dense point cloud completion. In *European Conference on Computer Vision*, pages 365–381. Springer, 2020.

[33] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.

[34] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *arXiv preprint arXiv:1905.10711*, 2019.

[35] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[36] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, pages 728–737. IEEE, 2018.